Visualizing Patterns of Drug Prescriptions with EventFlow: A Pilot Study of Asthma Medications in the Military Health System

Megan Monroe[1], Tamra E Meyer[2], PhD, Catherine Plaisant, PhD[1], Rongjian Lan[1], Krist Wongsuphasawat[1], PhD; Col. Trinka S Coster, MD[2], Sigfried Gold[3], Jeff Millstein[4], PhD, Ben Shneiderman, PhD[1]

[1]University of Maryland, College Park, MD;
[2]US Army, Office of the Surgeon General, Pharmacovigilance Center, Falls Church, VA;
[3] Social & Scientific Systems, Silver Spring, MD, [4] Oracle Health Sciences, Burlington, MA

Abstract

The Food and Drug Administration and Department of Defense were interested in detecting sub-optimal use of long-acting beta-agonists (LABAs) in asthmatics within the Military Health System (MHS). Visualizing the patterns of asthma medication use surrounding a LABA prescription is a quick way to detect possible sub-optimal use for further evaluation. The US Army, Office of the Surgeon General, Pharmacovigilance Center (PVC) selected a random sample of 100 asthma patients under age 65 with a new LABA prescription from January 1, 2006-March 1, 2010 in MHS healthcare claims. Analysis was conducted in EventFlow, a novel interactive visualization tool being developed by the University of Maryland Human Computer Interaction Lab (HCIL) to display and summarize time-point and interval data. EventFlow groups individuals that share the same sequence of medications and displays the average interval times between events. We found that EventFlow was effective in uncovering clinically relevant patterns in the data. Epidemiologists reported that EventFlow was a powerful tool for rapidly visualizing possible patterns of sub-optimal LABA use that can be targeted for intervention.

Introduction

Electronic Health Record (EHR) databases contain millions of patient records showing medication therapy (*e.g.*, initiation, change of dose, switch to a different medication, augmentation with additional medication for the same condition) for various medical conditions. Clinical researchers and quality improvement analysts have a need to characterize the nature and frequency of such patterns of medication use and, ultimately, to build and test hypotheses regarding the reasons for changes, the benefits and risks of real-world prescribing patterns, and to make recommendations for improved medication management. We believe that interactive information visualization tools such as EventFlow[1] (Figure 1) - developed by the University of Maryland Human Computer Interaction Lab (HCIL) - can lead to a dramatic improvement in a clinical researcher's ability to obtain deep knowledge of the types of medication usage patterns that occur within a medical practice/healthcare system.

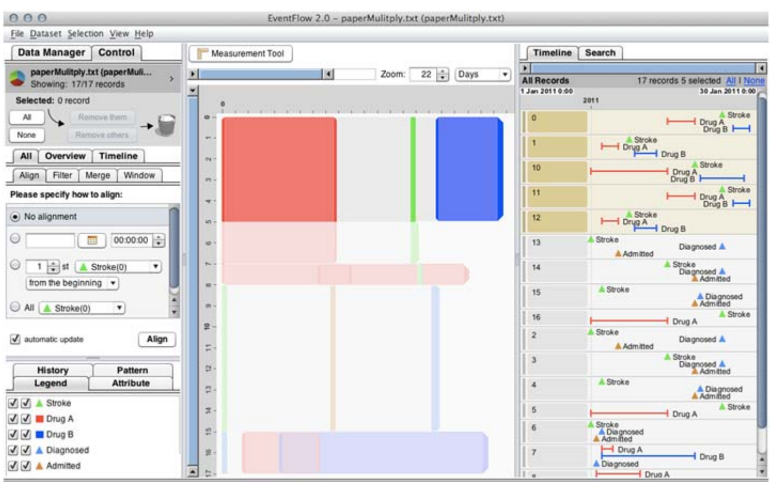

Figure 1: The EventFlow interactive analysis tool (www.cs.umd.edu/hcil/eventflow) with a small sample dataset. On the left are found controls and legend, in the middle is the overview of all sequence patterns in the dataset, and on the right a scrollable timeline browser shows all the individual records. The top sequence in the overview is selected (drug A, followed by stroke, followed by drug B). The distance between events corresponds to the average time between events. The height of the bar corresponds to the proportion of records with that sequence. The records with the selected sequence are highlighted at the top in the timeline view.

To evaluate novel approaches such as the one implemented in the EventFlow interface, designers and clinical researchers work together on Multifaceted In-Depth Longitudinal Case studies[2]. This paper reports on the early stages of a case study being conducted with the US Army, Office of the Surgeon General, Pharmacovigilance Center (PVC) where epidemiologists and clinicians worked with HCIL researchers to understand the prescribing patterns of asthma medications using our early prototypes. PVC is collaborating with the Food and Drug Administration (FDA)

Report Documentation Page

Form Approved
OMB No. 0704-0188

1. REPORT DATE **JUN 2013**	2. REPORT TYPE	3. DATES COVERED **00-00-2013 to 00-00-2013**

4. TITLE AND SUBTITLE **Visualizing Patterns of Drug Prescriptions with EventFlow: A Pilot Study of Asthma Medications in the Military Health System**	5a. CONTRACT NUMBER 5b. GRANT NUMBER 5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S)	5d. PROJECT NUMBER 5e. TASK NUMBER 5f. WORK UNIT NUMBER
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) **US Army, Office of the Surgeon General,Pharmacovigilance Center,5109 Leesburg Pike ,Falls Church,VA,22041-3258**	8. PERFORMING ORGANIZATION REPORT NUMBER
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSOR/MONITOR’S ACRONYM(S) 11. SPONSOR/MONITOR’S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release; distribution unlimited

13. SUPPLEMENTARY NOTES

14. ABSTRACT
The Food and Drug Administration and Department of Defense were interested in detecting sub-optimal use of long-acting beta-agonists (LABAs) in asthmatics within the Military Health System (MHS). Visualizing the patterns of asthma medication use surrounding a LABA prescription is a quick way to detect possible sub-optimal use for further evaluation. The US Army, Office of the Surgeon General, Pharmacovigilance Center (PVC) selected a random sample of 100 asthma patients under age 65 with a new LABA prescription from January 1, 2006-March 1 2010 in MHS healthcare claims. Analysis was conducted in EventFlow, a novel interactive visualization tool being developed by the University of Maryland Human Computer Interaction Lab (HCIL) to display and summarize timepoint and interval data. EventFlow groups individuals that share the same sequence of medications and displays the average interval times between events. We found that EventFlow was effective in uncovering clinically relevant patterns in the data. Epidemiologists reported that EventFlow was a powerful tool for rapidly visualizing possible patterns of sub-optimal LABA use that can be targeted for intervention.

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT **unclassified**	b. ABSTRACT **unclassified**	c. THIS PAGE **unclassified**	**Same as Report (SAR)**	**10**	

and Allergy/Immunology clinicians from the Water Reed National Military Medical Center to determine if long-acting beta-agonists (LABAs), a type of asthma medication, are being mis-prescribed within the Military Health System (MHS).

In this paper we briefly review previous work on temporal analysis tools and describe EventFlow, then introduce the general research questions which provide context for our pilot study, and finally report on the pilot study itself. We describe the process used to explore the data - illustrated with screen shots of the application - and summarize new insights and general feedback provided by users of Eventflow.

Temporal Analysis of Event Sequences - Related work

EventFlow builds on the past research by the HCIL on the interactive visualization of time series data, from early efforts on single patient history visualizations (LifeLines)[3] to our recent efforts on visualization of thousands of temporal sequences of point events (*i.e.*, events with a single timestamp) LifeLines2 and LifeFlow. With LifeLines2[4] researchers use event operators (align, rank, filter and group by) to specify queries on point event data. This early effort focused on searches for patterns specified by users, for example “find all the patients who bounce back to the ICU within 24 hours of leaving the ICU”, or “find patients with high creatinine readings within 14 days of the administration of radiographic contrast materials” (in an effort to find patients who experienced reduced renal function after infusion of contrast materials). Our research showed that such temporal queries require users to refine their queries iteratively after seeing the results (*e.g.*, by seeing the results they immediately realized that they needed to remove patients with too many creatinine highs before the contrast - as they probably had chronic renal failure - or that they needed to remove patients without normal readings before the contrast).

Our more recent project, LifeFlow[5], expands on LifeLines2 to provide compact visual summaries of all sequences found in the data. This breakthrough technique allows users to explore questions such as “what happens to patients after they leave the emergency room?” or, combined with the align operator “what happens before and after patients are admitted to the ICU?”. While those tools have been successfully used by clinical researchers and quality assurance administrators to answer many questions, they only operate on point event data (diagnoses, orders, admissions etc.) and therefore have no notion of episodes, partial or complete overlaps, or gaps between events.

EventFlow introduces the ability to interactively search and visualize interval data, which is an important step forward. Intervals, such as uninterrupted periods of medication use or episodes of disease, are a central aspect of the analysis of medication use.

In a broader context, temporal database storage, retrieval, interpretation, analysis, and visualization constitute a huge set of research topics that have been studied for a half century. While solutions for transactional systems such as financial trades have matured, many related temporal database problems remain. A key source of the difficulty is that temporal databases have a deceptively simple structure, usually relying on triples such as <patient ID#, event, timestamp>. While storage as a relational table is tempting, the self-joins to create a temporal history with hundreds or thousands of events make this simple approach inappropriate. In short, the relational database model can be too limiting for many temporal queries. Extensions such as TSQL and other temporal database query languages solve some problems, but the semantics of meaningful temporal queries is difficult or impossible to express in these temporal languages and logics, especially in the presence of interval events[6].

The complexity of EHR data (continuous and categorical) and the rich variety of questions that users seek to answer have stimulated an impressive assortment of solutions to temporal database problems. The first step is to extract appropriate subsets of data, clean the data to remove erroneous and redundant information, deal with missing and uncertain data, and then resolve domain specific problems such as gaps and overlaps in prescriptions[7]. We have found that visualization can vastly improve the efficiency of data cleaning and characterization efforts. For example, researchers can rapidly identify patterns that suggest data entry errors (*e.g.*, duplicate prescriptions) or can find drug usage that does not reflect administrative rules.

Data mining strategies that can extract common patterns in baskets of items, have been cleverly extended to deal with categorical sequences, that occur in many problems such as DNA analyses, sensor nets, and EHRs (*e.g.*, ref. 8-9). However, when events in a sequence have different amounts of time between them [*e.g.*, ref. 10], data mining approaches become difficult to apply, leading researchers to adopt more selective approaches that often include “interestingness” measures[11]. Work on extracting common temporal patterns in interval event data is still in its infancy[12].

EventFlow interface description

In this section we introduce the main design features of EventFlow. More details are provided in context in the later pilot study section. The main interface of EventFlow consists of three main components: interactive controls, legend, and timeline (Figure 1). The timeline browser shows details of individual records; each patient is shown on a

separate timeline (in Figure 1 we see patient 0 to 16, and users need to scroll to see all patients). Triangle icons represent point events (*e.g.*, green triangles represent stroke) while lines represent intervals (a red line represents drug A which becomes a small red rectangle when the interval is very short). The legend shows all event categories and users can change the color and order of the categories if desired. In the middle of the screen is the overview of all event sequences in the dataset. The aggregation principle was first introduced in[5]. All records with the same sequence of events are aggregated in a single bar. The height of a bar is determined by the number of records in the group and the horizontal gap between events is proportional to the mean time between the two events among the records in the group. Users can select other metrics such as the median, and the distribution of values is overlaid on the display when the cursor hovers over a time gap element. Multiple interval events can occur concurrently, and EventFlow handles this occurrence by rendering overlapping intervals using the combined color of the two overlapping categories. Colors selected for interval categories default to primary colors, resulting in intuitive overlap colors. For example, when a red interval intersects a blue interval, the resulting overlap is purple. When two intervals of the same category intersect, the color saturation in the overlapping region is increased. While this technique works when limited to a small number of event categories (***i.e.,*** colors) our experience has shown that being able to see overlaps of just two or three event categories was an important visualization effect for most users.

The two views (overview and timeline view) are coupled so that when users select an event sequence (*i.e.,* bar) all records with that sequence are selected on the timeline view (shown in dark yellow in Figure 1) and moved to the top of the timeline. Similarly, selecting a record on the right will highlight the corresponding event in the summary. The legend allows users to select or deselect which event categories they want to display on the overview or the timeline. After selecting records users can also remove either the selected or unselected records from the display. These two simple techniques allow users to easily narrow the focus of an analysis on records exhibiting particular event sequences of interest.

EventFlow also includes two separate search interfaces[21]. While the basic menu-based search interface gives users easy access to either before and after relationships (Subsequence module) or during relationships (Overlap module), Advanced Search allows users to specify these relationships in tandem, as well as access to more complex temporal features such as absolute time constraints and absence of events scenarios. The advanced search interface uses a visual query language to draw the desired sequence of event relationships (see Figure 7 of the pilot study section). When queries are specified, matching records are selected in the timeline display and moved to the top. This allows users to quickly see not only the records that were matched to their query, but also records that did not match so they can check that the search behaved as expected.

Finally the control panel gives access to many more powerful operators to zoom and filter, rank and cluster the records, adjust parameters of the views, and manage datasets.

Background: Ongoing PVC study of long-acting beta-agonist (LABA) therapy

FDA guidelines recommend that patients with asthma who are using LABA therapy also receive a long-term asthma controller medication (LTACM) such as an inhaled corticosteroid (ICS)[22-23]. Single-ingredient LABA prescribing has been linked with increased risk of asthma-related hospitalizations, intubations, and death. The FDA has issued drug safety communications and required label changes to all LABA products detailing that, in asthmatics, LABAs should only be prescribed in combination with a LTACM and only after attempts to optimize other LTACM therapies have failed. According to FDA guidelines LABA therapy should be de-escalated once asthma has been adequately controlled. Current LABA prescribing patterns in the MHS were unknown, so PVC and FDA undertook a study to characterize LABA use among asthmatics in the MHS over time. Metrics of interest were the proportion with concomitant LTACM therapies, length of therapy, and escalation/de-escalation practices. The ultimate goal of the study was to understand LABA use in the MHS in order to inform interventions to prevent morbidity and mortality from sub-optimal LABA use.

The PVC performed a retrospective cohort study using administrative health claims from the MHS from January 1, 2006 to March 1, 2010. The MHS provides care to over 9.5 million beneficiaries including active duty service members, retired service members, and their family members. The study was approved by the US Army Medical Research and Material Command, Office of Research Protections, Institutional Review Board. Data for this study were obtained from a data extract of the MHS Data Repository (MDR). The MDR is a compilation of data feeds of prescription records that supply information about dispense date, dose, quantity and days supply of a medication; medical encounter claims that supply information about diagnoses and medical procedures; and the Defense Eligibility and Enrollment Reporting System (DEERS) that provides information on MHS enrollment, eligibility for enrollment, and demographics. The data extract is updated on a quarterly basis and stored as an Oracle database.

Eligible patients were identified from the prescription files by searching for records containing the LABA product names “Symbicort”, “Advair”, “Dulera”, “Foradil”, “Performist”, or “Serevent” within the study period. Patients were further limited to those with asthma (identified from medical encounters using International Classification of Disease, 9th revision [ICD-9] code 493.xx) during the 365 days before the LABA dispensing and we selected their first incident LABA dispensing with no other LABA days supply in the previous 90 days. Patients ages 65 and older or those with codes for chronic obstructive pulmonary disorder (COPD; ICD-9 491.xx, 492.xx, or 496.xx) were further excluded in order to concentrate the evaluation on prescribing of LABAs among asthma patients. The LABA prescriptions considered in this study were only included if the patient was continuously eligible for enrollment in the MHS, was not outside the continental US, and was not deployed during the 365 days surrounding the LABA. These inclusion criteria improved the likelihood of capturing asthma patients while excluding COPD patients.

Case study of the use of EventFlow

From the 182,498 asthmatics under the age of 65 who had a new LABA prescription (index LABA) between January 2006 and March 2010 and met the other inclusion criteria, we randomly selected 100 patients, extracted all of their asthma medication prescriptions for the 365 days surrounding the index LABA prescription date, categorized the asthma medications into useful groups (LABA, ICS, short-acting beta-agonists [SABA], leukotriene

receptor antagonists [LTRA], mast stabilizers [MastStabil], oral corticosteroids [OCS], and older asthma therapies [Older]), and processed the file to import into EventFlow. The input file included a de-identified patient identifier, the asthma medication group identifier, the start date of the prescription, and the end date of the prescription derived from the start date and days supply (Figure 2). An indicator to mark the index LABA date was added to each patient's records. The indicator was needed because other LABA prescriptions could precede the index LABA prescription since only 90 days without a previous LABA was required to meet the index criteria of a new (incident) LABA. LABA prescriptions that occurred before the index LABA prescription were not eligible as index LABAs because they failed to meet one of the inclusion criteria.

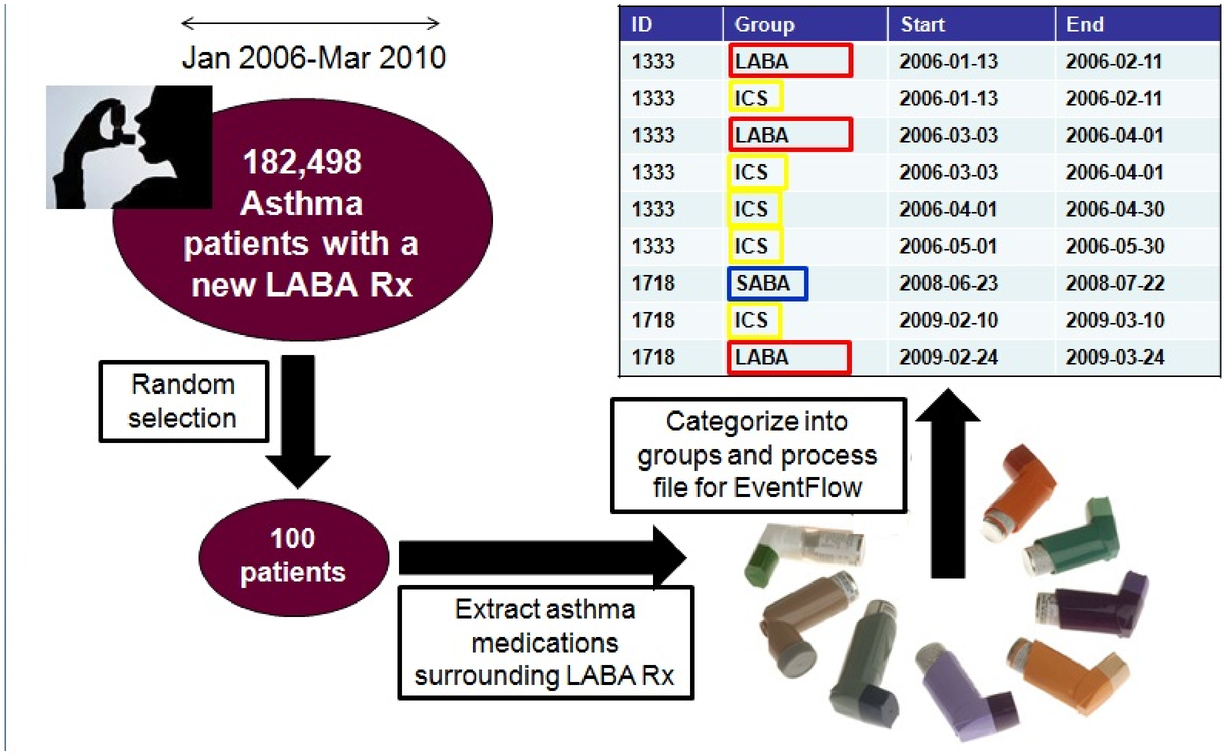

Figure 2: Data extraction process

We looked for patterns of asthma medication dispensing surrounding the index LABA prescription that suggested good compliance with FDA guidelines[22-23]. These patterns included 1) a LABA only in conjunction with an ICS, 2) use of other LTACM such as LTRA or ICS prior to the LABA, and 3) exchange of LABA for other LTACM after asthma is under good control.

The first goal was to gain an overview of the patterns of asthma prescriptions among the patients. We used the EventFlow visual overview of all asthma prescriptions found in the records (Figure 3). Among the seven possible asthma medication groups, the sample contained six of the groups. From the initial display, we noted that a large proportion of the patients' asthma therapy began with a LABA (red for LABA other than the index LABA prescription and green for the index LABA prescription). Starting asthma therapy with a LABA without first attempting control with other LTACM could be contrary to FDA guidelines so we examined patterns surrounding the index LABA further using the alignment tool in EventFlow.

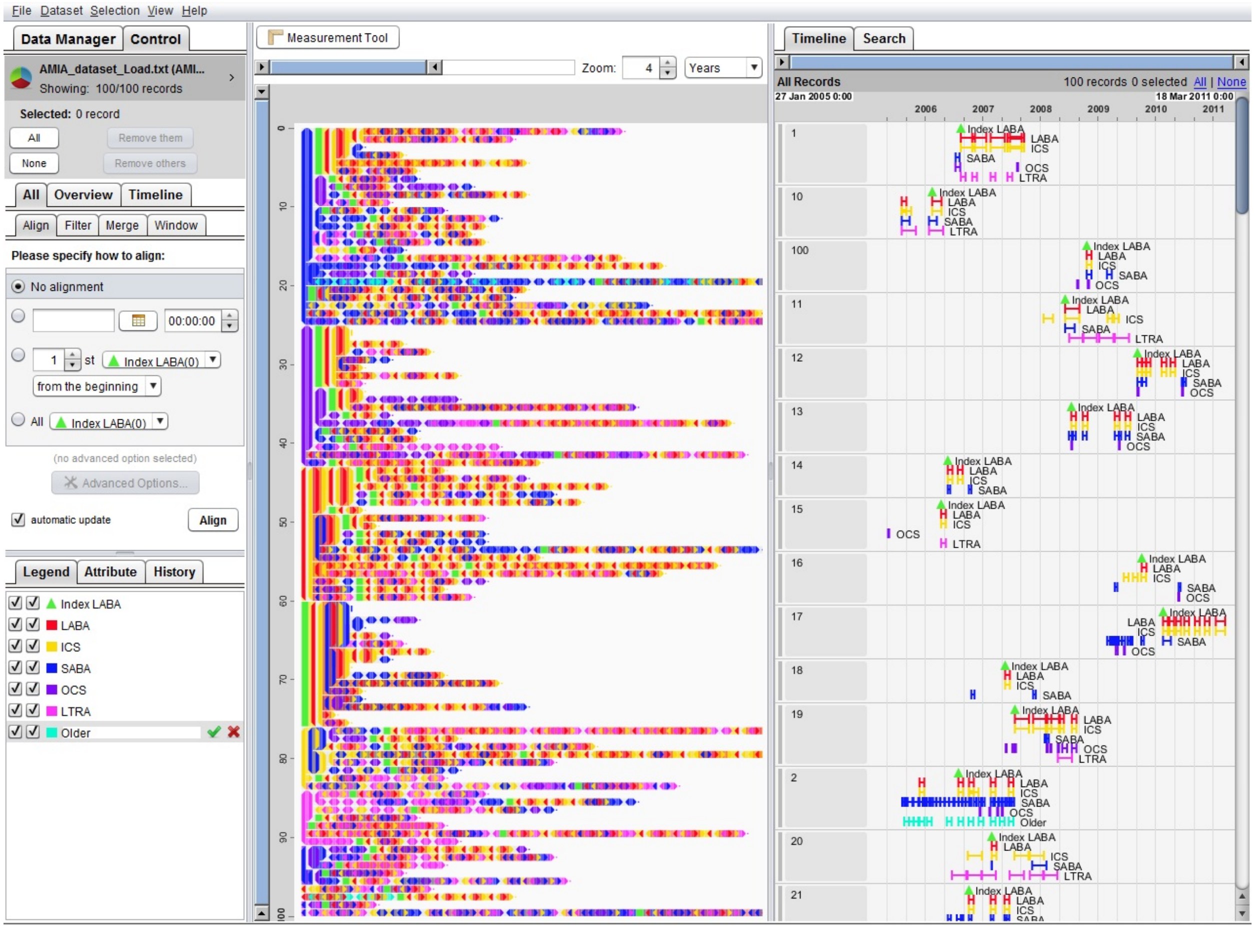

Figure 3: This screen shows EventFlow after the dataset of 100 patients' prescription records was loaded.

The next figure (Figure 4) shows the same data aligned on the point of interest, the index LABA prescription

From the aligned display, we noted that:

1. About 25% of this sample did not have evidence of other asthma medications before a LABA (*i.e.*,, started with red LABA or green index LABA). We were looking for other asthma medications like yellow ICS or pink LTRA before starting the LABA. However, starting on a LABA may be appropriate in the case of a significant asthma exacerbation or symptomatic presentation.
2. We noted that very few LABAs were prescribed without an inhaled corticosteroid (*i.e.*,, red LABA usually overlaps with the yellow ICS) which is consistent with FDA guidelines.
3. We saw some patients with long histories of blue SABA or purple OCS that may have benefited from earlier addition of a LABA.
4. Several patients had a non-index LABA greater than 90 days but less than a year before the index LABA suggesting that a longer washout period may be useful in identifying new (incident) LABA.

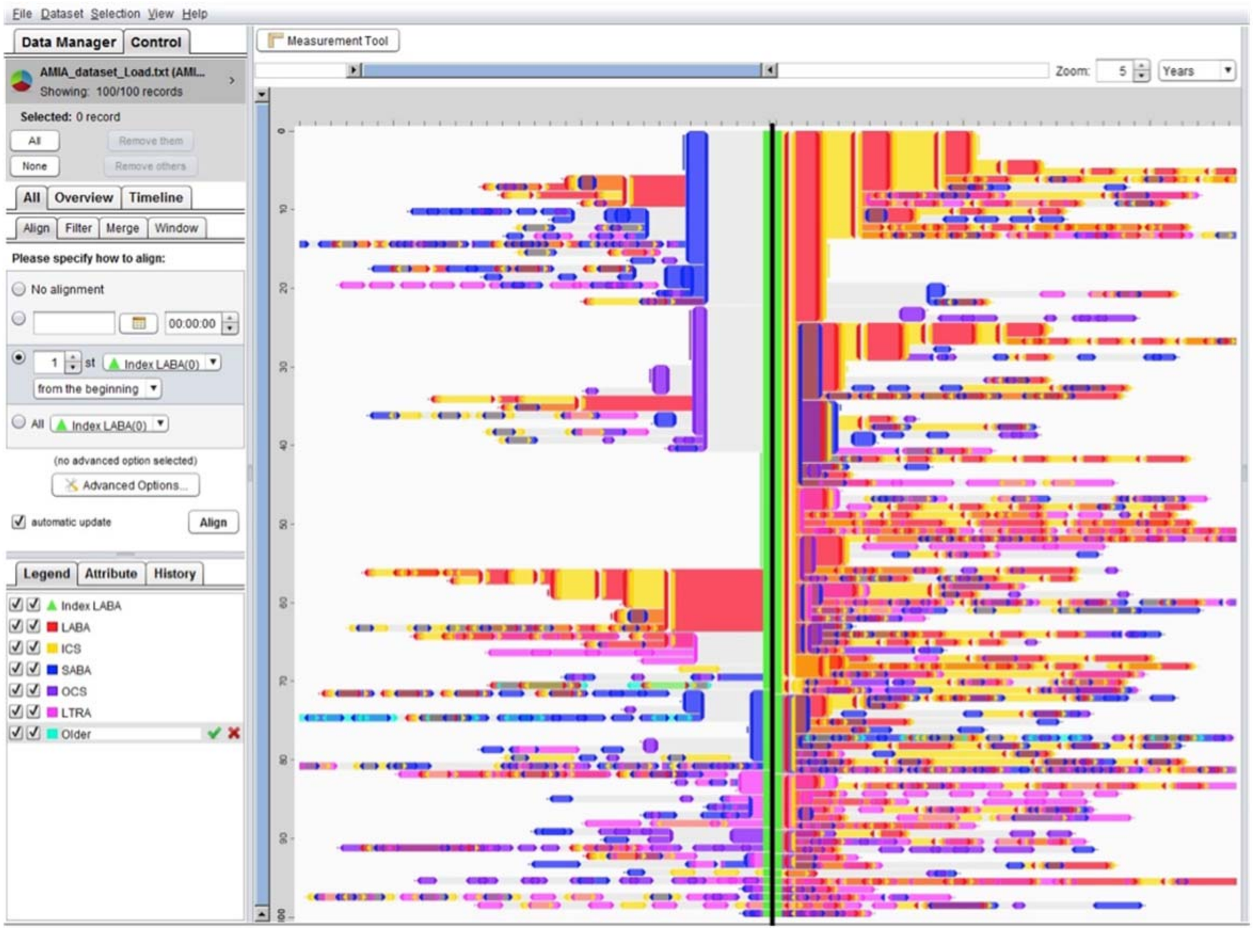

Figure 4: The same data is now aligned on the green index LABA (the alignment point is shown as a black vertical line). Note that Eventflow generates 2 separate overviews: one for the period preceding the index LABA prescription (on the left of the vertical black line) and one for the period following the index LABA prescription (on the right of the black line), therefore elements on the left do NOT match elements on the right at the same vertical position.

To focus on patterns of step-up therapy from a non-LABA LTACM to a LABA and step-down therapy from a LABA to other LTACM, we limited the display to LABA and the other common LTACM: ICS and LTRA (Figure 5). In the absence of significant asthma exacerbations or clinical presentations, we expected to see step up from an ICS or LTRA to a LABA and then step down to an LTRA or ICS alone once asthma was under good control. However, we saw that about 65% of patients did not have another LTACM before the LABA and about half had no evidence of step down to an LTRA or ICS alone in this pilot cohort. Again, these patterns are not necessarily inconsistent with FDA guidelines[22-23] if asthma symptoms warranted immediate or continued use of LABA. Nevertheless, these patterns do highlight areas for further investigation and possible clinical intervention of those patients with no evidence of a step-down trial. We are currently working with allergy/immunology specialists at Walter Reed National Military Medical Center, Bethesda, Maryland to evaluate patients who continued on LABA therapy but lacked clinical outpatient follow-up visits after the index LABA.

Finally, we constructed multiple queries to view the patterns that were consistent with FDA guidelines[22-23] by finding patients who were prescribed an ICS, the most common LTACM, within 3 months of both the start and end date of a LABA prescription. This query returned some patients who had an ICS within 3 months of both the start and end of one of their LABA prescriptions yet still started or ended therapy on a LABA contrary to FDA guidelines. Therefore, we added the limitation that the patient therapy should not start or end with a LABA (Figure 6). The query identified five patients that met the FDA guidelines based on prescription patterns of step-up therapy from an ICS to a LABA, use of LABA concomitantly with an ICS, and step-down to an ICS alone.

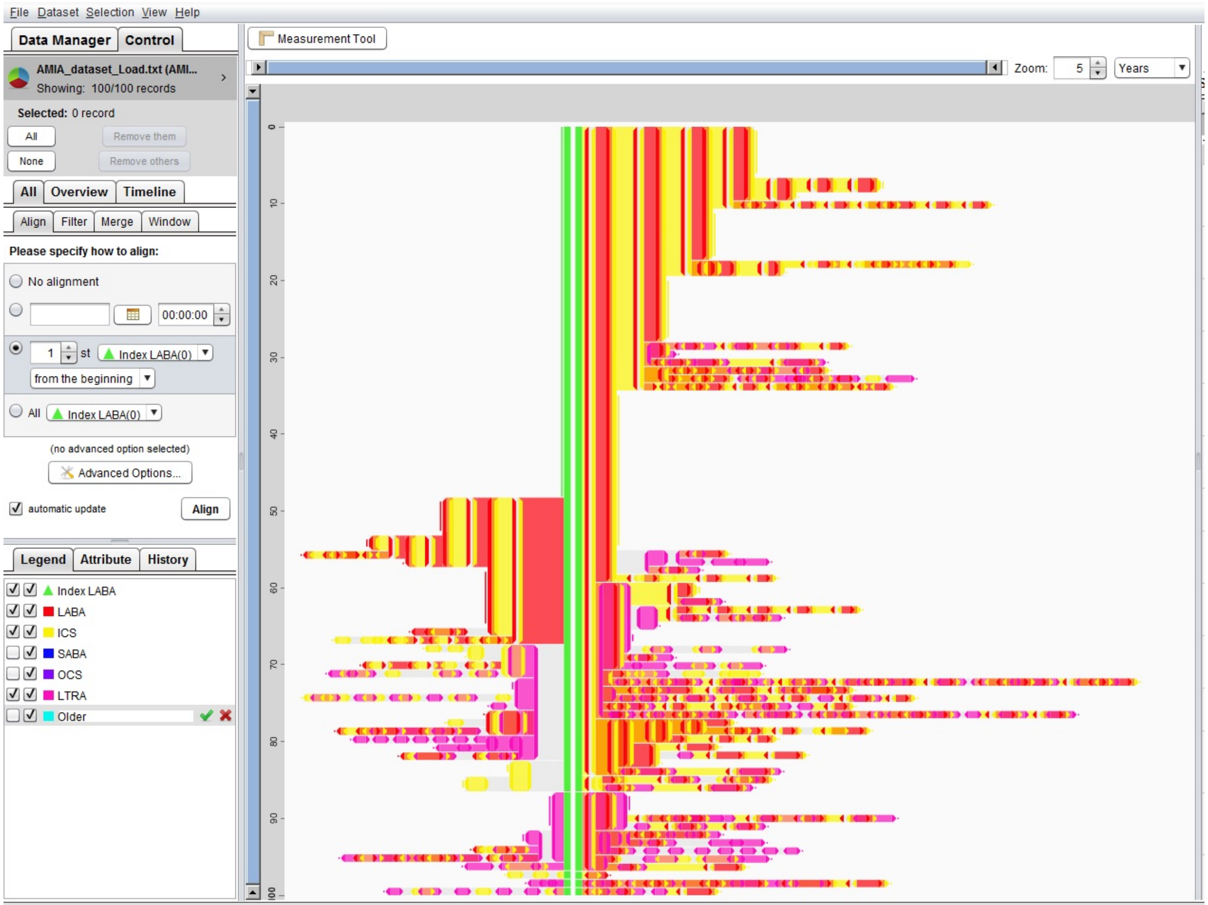

Figure 5: Limiting the display to LABA, ICS, and LTRA simplifies the display.

Because the exploration was done on a sample of 100 patients, to confirm these exploratory findings, the entire patient population was evaluated using traditional SQL and SAS queries. These analyses supported the exploratory results.

The epidemiologists and clinicians reported that both query interfaces were easy to use. They commented that it was much easier to learn and use than the command-based, statistical software that they normally employ. Furthermore, these command-based languages offered no simple option of visualizing the patterns in a meaningful way - a feature that is central to the standard search model as well as both of our query interfaces.

They noted that the query tools and visualization of results allowed them to rapidly develop hypotheses and test them “on the fly” without additional coding thereby supporting the general hypothesis that well-designed interactive visualization facilitates analysis of drug utilization practices.

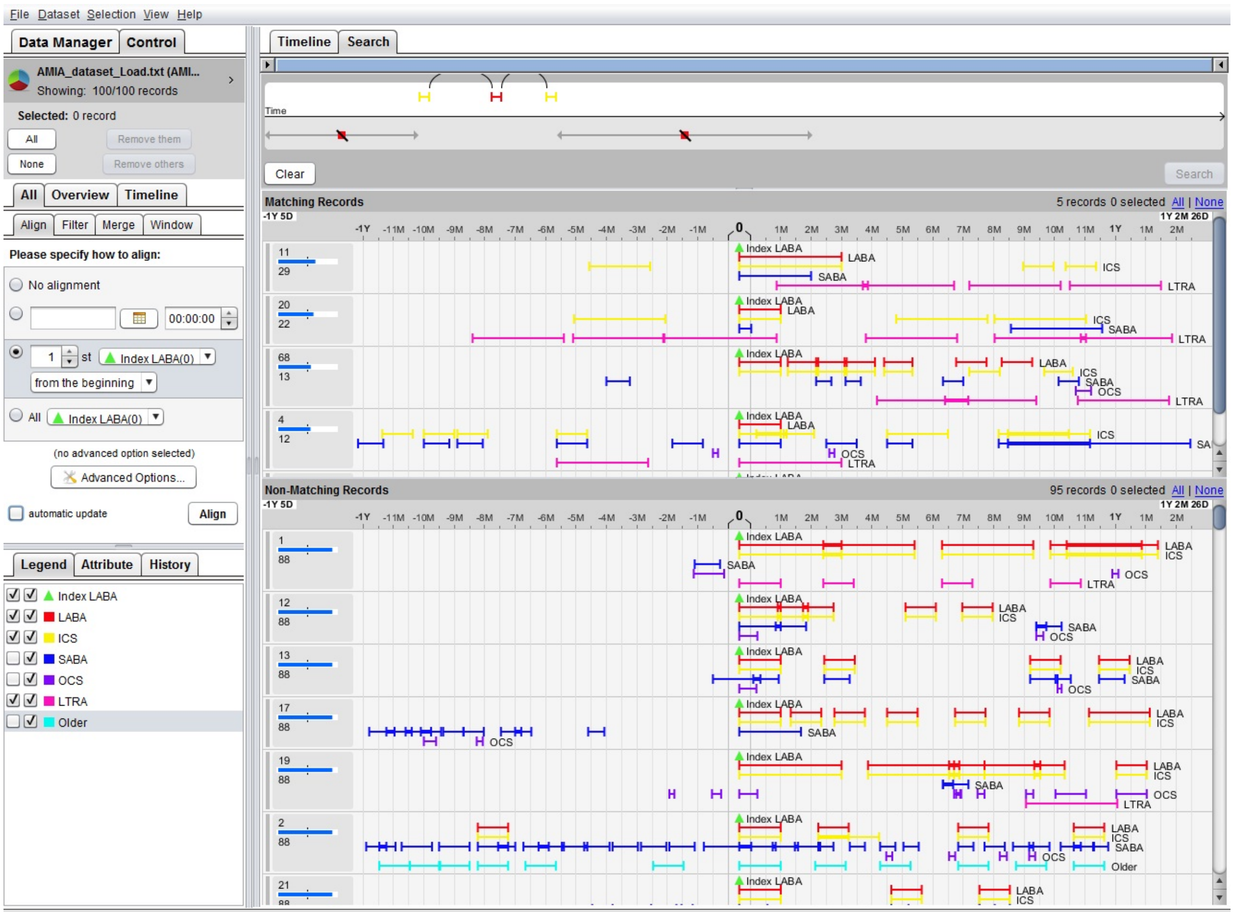

Figure 6: Using the advanced query panel, we searched for patients who received an ICS (in yellow) within 3 months of both the start and end date of the LABA (in red). We also wanted to ensure that this sequence was neither preceded nor followed by a LABA

Conclusion

EventFlow was effective in uncovering important patterns of LABA use in the data used in this pilot study. Our results provide evidence that secondary use of electronic health records databases could be dramatically expanded if easy-to-use interfaces allowed clinical researchers to specify queries, review results, and find temporal patterns without needing to default to the time-consuming traditional SQL or SAS queries. We believe that the future of user interfaces is moving toward larger, information- abundant interactive visual displays similar to EventFlow, and this will help researchers compare populations, discover relationships, and spot anomalies that are medically-actionable.

Acknowledgements

The University of Maryland appreciates the partial support of the Oracle Corporation for this research. Funds for the LABA prescribing practices study were received from FDA Safe Use Initiative. We wish to thank Drs. Cecilia Mikita and Maureen Petersen, clinicians with Walter Reed National Military Medical Center, Bethesda, Maryland for their input into pattern recognition of sub-optimal therapy, and we also thank David Wang for his major contributions to our early work (as part of his PhD Thesis work on LifeLines 2).

References

1. EventFlow Webpage: www.cs.umd.edu/hcil/eventflow
2. Shneiderman, B. and Plaisant, C., Strategies for evaluating information visualization tools: multi-dimensional in-depth long-term case studies Proceedings of the 2006 AVI workshop on BEyond time and errors: novel evaluation methods for information visualization (2006) 1-7

3. Plaisant, C., Mushlin, R., Snyder, A., Li, J., Heller, D., and Shneiderman, B., LifeLines: Using visualization to enhance navigation and analysis of patient records, American Medical Informatics Association 1998 Annual Fall Symposium (Orlando, Nov. 9-11, 1998) AMIA, Bethesda MD, 76-80 (1998).
4. Wang, T. D., Wongsuphasawat, K., Plaisant, C., and Shneiderman, B., Extracting insights from Electronic Health Records: Case studies, a visual analytics process model, and design recommendations, *Journal of Medical Systems* 35, 5 (2011), 1135-1152.
5. Wongsuphasawat, K., Gomez, J. A. G., Plaisant, C., Wang, T. D., Shneiderman, B., and Taieb-Maimon, M., LifeFlow: Visualizing an overview of event sequences, *Proc. ACM SIGCHI Conference*, ACM Press, New York (2011), 1747-1756
6. Jensen C. S. and R. T. Snodgrass, Temporal data management, IEEE Transactions on Knowledge and Data Engineering, vol. 11 (1999), 36-44.
7. Kandel, S., Heer, J., Plaisant, C., Kennedy, J., an Ham, F., Riche, N. H., Weaver, C., Lee, B., Brodbeck, D., and Buono, P., Research directions in data wrangling: Visualizations and transformations for usable and credible data, *Information Visualization 10* (October 2011), 271-288.
8. Garofalakis, M., Rastogi, R. and Shim K. Mining sequential patterns with regular expression constraints, *IEEE Transactions on Knowledge and Data Engineering, 14,* 3 (2002), 530–552.
9. de Amo, S. and Furtado D. A.. First-order temporal pattern mining with regular expression constraints. Data & Knowledge Engineering, 62(3):401–420, 2007.
10. Zaki, M. J., Sequence mining in categorical domains: Incorporating constraints, Proc. Conf. on Information and Knowledge Management (2001), 422-429.
11. Lenca, P., Meyer, P., Vaillant,B. and Lallich, S., On selecting interestingness measures for association rules: User oriented description and multiple criteria decision aid, *European Journal of Operational Research* 184, 2 (2008), 610-626,.
12. Patel, D., Hsu, W., and Lee, M., Mining Relationships Among Interval-based Events for Classification, Proc. of the 2008 ACM SIGMOD Conference, 2008, 393-404.
13. Keim, D., Kohlhammer, J., Ellis, G., and Mansmann, G. (Editors), Mastering the Information Age: Solving Problems with Visual Analytics, Eurographics Association, Goslar, Germany (2010).
14. Noren, G. N., Bate, A., Hopstadius, J., Star, K. and Edwards, I. R., Temporal pattern discovery for trends and transient effects: its application to patient records, Proc. 14th International Conference on Knowledge Discovery and Data Mining (2008), 963-971.
15. Murphy, S, M. Mendis, K. Hackett, R. Kuttan,W. Pan, L. Phillips, V. Gainer, D. Berkowicz, J. Glaser, I. Kohane, and H. Chueh. Architecture of the opensource clinical research chart from informatics for integrating biology and the bedside. Proceedings of AMIA (2007).
16. Post, A. R.and Harrison, J. H., Protempa: A method for specifying and identifying temporal sequences in retrospective data for patient selection. *Journal of the American Medical Informatics Association* (2007).
17. Stacey, M., McGregor, C., Temporal abstraction in intelligent clinical data analysis: A survey, Artificial Intelligence in Medicine (2007) 39, 1—24
18. Aigner, W., Miksch, S., Schumann, H. and Tominski, C., *Visualization of Time-Oriented Data*. Springer, Berlin (2011).
19. A. Rind, T. D. Wang, W. Aigner, S. Miksch, K. Wongsuphasawat, C. Plaisant, and B. Shneiderman, Interactive Information Visualization to Explore and Query Electronic Health Records, *Foundations and Trends in Human–Computer Interaction*, Vol. 5, No. 3 (2013) 207–298
20. Jin, J. and Szekely, P., Interactive querying of temporal data using a comic strip metaphor, Proc. IEEE Symposium on Visual Analytics Science and Technology (VAST) (2010).
21. Monroe, M., Lan, R., Morales del Olmo, J., Shneiderman, B., Plaisant, C., Millstein, J., The Challenges of Specifying Intervals and Absences in Temporal Queries: A Graphical Language Approach, *Proc. of ACM Conference on Human-Computer Interaction,* (2013) to appear
22. FDA 2010a: Food and Drug Administration. FDA drug safety communication: new safety requirements for long-acting inhaled asthma medications called Long-Acting Beta-Agonists (LABAs). 2010; http://www.fda.gov/Drugs/DrugSafety/PostmarketDrugSafetyInformationforPatientsandProviders/ucm200776.htm. Accessed June 28, 2011.
23. FDA 2010b: Food and Drug Administration. FDA drug safety communication: drug labels now contain updated recommendation on the appropriate use of long-acting inhaled asthma medications called Long-Acting Beta-Agonists (LABAs). 2010; http://www.fda.gov/Drugs/DrugSafety/PostmarketDrugSafetyInformationforPatientsandProviders/ucm213836.htm Accessed June 28, 2011.